manuscript No.
(will be inserted by the editor)

State dependent swap strategies and adaptive adjusting of number of temperatures in Parallel Tempering algorithms.

Mateusz Krzysztof Łącki · Błażej Miasojedow

arXiv:1406.3218v2 [stat.CO] 24 Jun 2014

Received: date / Accepted: date

Abstract In this paper we present extensions to the original adaptive *Parallel Tempering* algorithm. Two different approaches are presented. In the first one we introduce state-dependent strategies using current information to perform a swap step. It encompasses a wide family of potential moves including the standard one and *Equi Energy* type move, without any loss in tractability. In the second one, we introduce online adjustment of the number of temperatures. Numerical experiments demonstrate the effectiveness of the proposed method.

Keywords Parallel Tempering · Adaptive MCMC · Swapping Strategies · Equi-Energy Sampler

1 Introduction

Markov chain Monte Carlo (MCMC) is a generic method to approximate an integral of the form

$$I := \int_{\mathbb{R}^d} f(y)\pi(y)\mathrm{d}y \ ,$$

where π is a probability density function, which can be evaluated point-wise up to a normalising constant. Such an integral occurs frequently when computing Bayesian posterior expectations (Robert and Casella, 1999; Gilks et al, 1998).

The random-walk Metropolis algorithm (Metropolis et al, 1953) often works well, provided the target density π is, roughly speaking, sufficiently close to unimodal. The efficiency of the Metropolis algorithm can be optimised by a suitable choice of proposal distribution. These, in turn, can be chosen automatically by several adaptive MCMC algorithms; see Haario et al (2001); Atchadé and Rosenthal (2005); Roberts and Rosenthal (2009); Andrieu and Thoms (2008) and references therein.

When π has multiple well-separated modes, the random-walk based methods tend to stuck in a single mode for long periods of time. It can lead to false convergence and severely erroneous results. Using a tailored Metropolis-Hastings algorithm can help, but, in many cases, finding a good proposal distribution is not easy. Tempering of π, that is, considering auxiliary distributions with density proportional to π^β with $\beta \in (0,1)$, often provides better mixing between modes (Swendsen and Wang, 1986; Marinari and Parisi, 1992; Hansmann, 1997; Woodard et al, 2009).

We focus here particularly on the parallel tempering algorithm, which is also known as the replica exchange Monte Carlo and the Metropolis coupled Markov chain Monte Carlo.

The tempering approach is particularly tempting in such settings where π admits a physical interpretation, and there is good intuition how to choose the temperature schedule for the algorithm.

In general, choosing the temperature schedule is a non-trivial task, but there are generic guidelines for temperature selection based on both empirical findings and theoretical analysis (Kofke, 2002; Kone and Kofke,

Work partially suported by Polish National Science Center grant No. N N201 608740 and grant No. 2011/01/B/NZ2/00864

B. Miasojedow
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
Tel.: +48-22-55-44-441
E-mail: B.Miasojedow@mimuw.edu.pl

M.K. Łącki
Institute of Informatics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland
Tel.: +48-602-39-50-10
E-mail: Mateusz.Lacki@biol.uw.edu.pl

2005; Atchadé et al, 2011; Roberts and Rosenthal, 2012). These theoretical findings were used to derive adaptive version of the *Parallel Tempering* (Miasojedow et al, 2013a).

In the present paper we consider the adaptive version of the Parallel Tempering algorithm. The adaption consists in introducing state dependent swaps between differently tempered random walks. We study the impact of different distributions on potential steps and call them *Strategies*. Our choice of strategies is driven by solutions already known to the literature (Kou et al, 2006) and used within *Parallel Tempering* algorithm by Baragatti et al (2013). The novelty of our approach stems from an alternative implementation of *Equi Energy* moves that renders the algorithm parameters free, i.e. the user does not need to provide precise *Energy Rings* any more. We also investigate different modifications of this new approach.

We also propose an automated method for adapting the actual number of considered temperatures, in the spirit of Miasojedow et al (2013a). The temperature adaptation scheme depends on the parameters of the adaptive random walks applied in the parallelised Metropolis-Hastings stage of the algorithm in case when the state space amounts to be the usual $\mathbb{R}^d$.

We have also showed that the proposed algorithm satisfies the Law of Large numbers, in the same setting as in Miasojedow et al (2013a).

2 Definition and Notations

Our basic object of interest is the density $\pi : \Omega \mapsto \mathbb{R}_+$, where $\Omega = \mathbb{R}^d$. We assume we can evaluate pointwise a function that is proportional to π by some constant. The *Parallel Tempering* approach suggests to construct a Markov chain on the product space Ω^L, where L is the number of temperature levels. On that space a new density π^β is constructed by posing for $x \in \Omega^L$

$$\pi^\beta(x) = \pi^\beta(x_1, \dots, x_L) \propto \pi^{\beta_1}(x_1) \times \cdots \times \pi^{\beta_L}(x_L)$$

so that $x_i \in \Omega$, $\beta = \left(\beta_1, \beta_2, \dots, \beta_L\right)$ are the inverse temperatures subject to $1 = \beta_1 > \beta_2 > \cdots > \beta_L$.

Vector $T = (T_1, \dots, T_L)$, where $T_\ell = \beta_\ell^{-1}$, contains numbers known as temperatures and is itself referred to as the temperature scheme.

Density π^β is known up to proportionality factor and by marginalising it w.r. to the first coordinate we retrieve the original distribution π.

Markov chain $X = \{X^{[k]}\}_{k \geq 0}$ targets π^β and can be spit into L coordinate chains, since $X^{[k]} = \left(X_1^{[k]}, \dots X_L^{[k]}\right)$. First coordinate chain will be referred to as the base chain.

The main idea behind *Parallel Tempering* is to interweave random walk steps with random swaps between chains. Each random swap exchanges results of a random walk step from two coordinate chains. Chains corresponding to higher temperatures [1] should in principle be more volatile and travel between different modes more easily than chains linked to lower temperatures. For if x is the last visited place by the l^{th} chain, and y is a proposal drawn from a region where the density assumes smaller values, $\pi(y) < \pi(x)$, then the probability of accepting such proposal, equal to η_ℓ, is higher on the more tempered chain[2]

$$\eta_\ell(x, y) = 1 \wedge \left(\frac{\pi(y)}{\pi(x)}\right)^{\beta_l} > 1 \wedge \frac{\pi(y)}{\pi(x)} = \eta_1(x, y).$$

Therefore, the more exchanges of higher tempered chains with the base chain, the bigger the chance of escaping quickly local probability clusters where a simple Markov chain would usually get stuck.

Formally, the generation of Markov chain X proceeds by repeated application of two kernels,

$$X^{[n-1]} \xrightarrow{\mathcal{M}} \widetilde{X}^{[n-1]} \xrightarrow{\mathcal{S}} X^{[n]}.$$

The random walk kernel $\mathcal{M}$ is defined on the product sigma algebra generator by

$$\mathcal{M}(x, A_1 \times \cdots \times A_L) = \prod_{l=1}^{L} \mathcal{M}_l(x_l, A_l),$$

where

$$\mathcal{M}_l(x_l, A_l) \equiv \int_A \eta_\ell(x_l, y_l) q(y_l - x_l) \mathrm{d}\, y_l + \mathbf{1}_{\{x_l \in A\}} \mathrm{r}(x_l),$$

and corresponds to independently performing random walk on different chains.

Let T_{ij} denote the transposition of coordinate i and j of

$$T_{ij}(x) = (x_1, \dots, x_j, \dots, x_i, \dots, x_L).$$

Then the proposal measure for the random swap is discrete and potentially state-dependent and is given by

$$\mathcal{Q}(x, A) \equiv \sum_{i<j} p_{ij}(x) \mathbf{1}_{\{A\}}(T_{ij} x), \tag{1}$$

where we one is free to choose the precise relation between p_{ij} and x, establishing with what probabilities should one carry out the transposition.

[1] Chains with lower β.

[2] Here we assume that random walk proposal kernel is symmetric.

The transition kernel of the swap step is defined simply by

$$\mathcal{S}(x,A) := \sum_{i<j} \mathbf{1}_{\{T_{ij}(x)\in A\}} p_{ij}(x)\alpha_{ij}(x) + \mathbf{1}_{\{x\in A\}} \sum_{i<j} \Big(1-\alpha_{ij}(x)\Big) p_{ij}(x)\ . \tag{2}$$

The precise form of the acceptance probabilities will be presented in next section.

3 State dependent swap strategies

In standard Parallel Tempering algorithm the proposal of swap step is drawn from distribution independent on the current state of process. Usual choice is to sample pair of temperatures uniformly at random or propose only pair of adjacent temperatures chosen also uniformly. Motivation of studies strategies of swap dependent on current state is Equi-Energy sampler proposed by Kou et al (2006) and adopted to *Parallel Tempering* algorithm in Baragatti et al (2013).

The main idea behind these algorithms is to exchange states with similar energy (i.e. value of log-density). The original *Equi Energy* sampler is greedy in memory usage: it must store all points drawn from differently tempered trajectories. This might become problematic in high dimensional settings. *Equi Energy* also targets a biased distribution instead of the proper one due to the use of asymptotic values in acceptance probability. Running this algorithm also requires specifying in advance the so called Energy Rings, i.e. the partition of the state space into regions with similar energy. Schreck et al (2013) addresses the problem of choosing Energy Rings and provides an and adaptive version of algorithm.

The Parallel Tempering algorithm with Equi-Energy moves proposed by Baragatti et al (2013) also requires to specify Energy-Rings. This hindrance can be circumvented with the use of state dependent swap strategies performing Equi-Energy-like moves inside the Parallel Tempering algorithm with no need to predefine Energy Rings themselves. In addition, our approach is flexible and using different strategies one can promote large jumps or other features.

The general algorithm is as follows: given current state of process after the random walk phase, $x = (x_1, \ldots, x_L)$, we propose that the drawing of transposition[3] $T_{ij}(x)$ should follow the state dependent discrete distribution defined by some probabilities $p_{ij}(x)$, defined on a discrete simplex $i < j$.

[3] We exchangeably refer to coordinate transposition as swaps.

To assure reversibility, the swap acceptance probability should be defined by

$$\alpha_{ij}(x) := \frac{p_{ij}\big(T_{ij}(x)\big)}{p_{ij}(x)} \left(\frac{\pi(x_i)}{\pi(x_j)}\right)^{\beta_j-\beta_i} \wedge 1\ . \tag{3}$$

The definition of acceptance probability (3) assures that kernel $\mathcal{S}$ is reversible with respect to $\boldsymbol{\pi}$.

Proposition 1 *Kernel $\mathcal{S}$ defined by (2) is reversible with respect to $\boldsymbol{\pi}(x) \propto \pi(x_1)^{\beta_1} \times \cdots \times \pi(x_L)^{\beta_L}$.*

Proof We need to show that for all $A, B \in \mathfrak{F}$ we have

$$\int_{A\times B} \boldsymbol{\pi}(\mathrm{d}x)\mathcal{S}(x,\mathrm{d}y) = \int_{A\times B} \boldsymbol{\pi}(\mathrm{d}y)\mathcal{S}(y,\mathrm{d}x)\ .$$

For all $i < j$ let us define

$$\mathcal{S}_{ij}(x,A) := \mathbf{1}_{\{T_{ij}(x)\in A\}} p_{ij}(x)\alpha_{ij}(x)\ .$$

It is enough to verify that for every $i < j$

$$\int_{A\times B} \boldsymbol{\pi}(\mathrm{d}x)\mathcal{S}_{ij}(x,\mathrm{d}y) = \int_{A\times B} \boldsymbol{\pi}(\mathrm{d}y)\mathcal{S}_{ij}(y,\mathrm{d}x)\ .$$

For any arbitrary chosen $i < j$ define a measure μ on $\mathbb{R}^{2dL}$ as follows: for $A, B \in \mathfrak{F}$ let

$$\mu(A\times B) := \lambda_{\mathrm{Leb}}\Big(\big\{x \in A\ :\ T_{ij}(x) \in B\big)\big\}\Big)\ ,$$

where λ_{Leb} denotes the Lebesgue measure on $\mathbb{R}^{dL}$.

Since $T_{ij}\big(T_{ij}(x)\big) = x$, by symmetry of Lebesgue measure we get

$$\begin{aligned}\mu(A\times B) &= \lambda_{\mathrm{Leb}}\Big(A \cap T_{ij}(B)\Big) = \\ &= \lambda_{\mathrm{Leb}}\Big(T_{ij}(A) \cap B\Big) = \mu(B\times ß)\ ,\end{aligned} \tag{4}$$

and by definition of $\mathcal{S}_{ij}$ we obtain that

$$\begin{aligned}\int_{A\times B} &\boldsymbol{\pi}(\mathrm{d}x)\mathcal{S}_{ij}(x,\mathrm{d}y) = \\ &= \int_{A\times B} \boldsymbol{\pi}(x)\alpha_{ij}(x)p_{ij}(x)\mu(\mathrm{d}x,\mathrm{d}y) = \\ &= \int_{A\cap T_{ij}(B)} \boldsymbol{\pi}(x)\alpha_{ij}(x)p_{ij}(x)\mathrm{d}x\end{aligned} \tag{5}$$

Now, using (3) we find that

$$\boldsymbol{\pi}(x)\alpha_{ij}(x)p_{ij}(x) = \boldsymbol{\pi}\big(T_{ij}(x)\big)\alpha_{ij}\big(T_{ij}(x)\big)p_{ij}\big(T_{ij}(x)\big)\ ,$$

and setting $y = T_{ij}(x)$ to (5) and applying (4) we get

$$\begin{aligned}\int_{A\times B} &\boldsymbol{\pi}(\mathrm{d}x)\mathcal{S}_{ij}(x,\mathrm{d}y) = \\ &= \int_{T_{ij}(A)\cap B} \boldsymbol{\pi}(x)\alpha_{ij}(x)p_{ij}(x)\mathrm{d}x = \\ &= \int_{A\times B} \boldsymbol{\pi}(x)\alpha_{ij}(x)p_{ij}(x)\mu(\mathrm{d}y,\mathrm{d}x) = \\ &= \int_{A\times B} \boldsymbol{\pi}(\mathrm{d}y)\mathcal{S}_{ij}(y,\mathrm{d}x),\end{aligned}$$

which completes the proof.

Remark 1 Thanks to the definition of kernel $\mathcal{S}$, for any positive measurable function $F : \mathbb{R}^{Ld} \to \mathbb{R}^+$ invariant by permutation we get $\mathcal{S}F(x) = F(x)$ which, under some regularity conditions (Miasojedow et al, 2013a), implies that *Parallel Tempering* algorithm with state dependent swap steps is geometrically ergodic. For under the same assumptions THEOREM 1 precised in the above-mentioned *œuvre* holds with state dependent swap steps.

To pursue an *Equi Energy* type moves without a need to fine-tune some additional parameters, e.g. precising *Energy Rings*, we propose to set the swap probabilities as follows: for $i < j$ let

$$p_{ij}(x) \propto \exp\big\{ -|\log(\pi(x_i) - \log(\pi(x_j)|\big\} . \tag{6}$$

The norming constant is equal $\sum_{i<j} p_{ij}(x)$, so it does not depend upon permutation of x, thus $p_{ij}(x) = p_{ij}\big(T_{ij}(x)\big)$, and further acceptance probability simplifies to acceptance probability of standard Parallel Tempering algorithm

$$\alpha_{ij}(x) = \left(\frac{\pi(x_i)}{\pi(x_j)}\right)^{\beta_j - \beta_i} \wedge 1 . \tag{7}$$

In comparison with standard PT, the state dependent swaps version propose more often states which are close in the sense of difference of energy. But, since acceptance mechanism is the same this leads to increasing number of global moves in the algorithm. The simulations results, presented below, confirm this improvement.

Other possible strategies commonly used in the literature include swapping at random (RA) all possible pairs, thus setting $p_{ij} = \binom{L}{2}^{-1}$, and swapping at random only the adjacent levels (AL), where $p_{ij} = (L-1)^{-1}\mathbf{1}_{\{i=j-1\}}$.

Remark 2 Note, that PT with Equi-Energy moves can be considered a special case of state dependent swap step. Let $H_1, \ldots, H_m$ be the Energy Rings. Denote by H_x the set H_i where x belongs to. Setting $p_{ij}(x) \propto \mathbf{1}_{\{x_j \in H_{x_i}\}}$, the acceptance probabilities also reduce to (7). In the end we obtain an algorithm proposed by Baragatti et al (2013).

Therefore the theoretical results presented in the present paper can be directly applied to PTEEM. In particular, it is geometrically ergodic under the same regularity condition, also convergence and the Law of Large Numbers for PTEEM with adaptive Metropolis step at each level can be thus easily obtained.

4 Adjusting number of temperature levels

The choice of parameters, i.e. the schedule of temperatures and parameters of random walk Metropolis at each level, is crucial for performance of the PT algorithm. This becomes apparent especially in case of the temperatures not assuming direct physical interpretation.

In a recent paper Miasojedow et al (2013a) has proposed an adaptive scheme to tune both temperature schedule and proposal distributions of the random walk steps at each temperature level. The above-mentioned algorithm still requires user-provided number of temperature levels and usually some pilot runs of the algorithm are necessary to determine their number.

In this section a simple criterion is presented for probing whether the algorithm has attained the correct number of temperature levels during its run-time. An incorrect temperature schedule leads to unsatisfactory estimation results: underestimating the number of temperature levels might lead the most-tempered chain to explore a distribution that is still multimodal, contradicting the whole idea of parallel tempering. On the other hand, overestimating levels number might lead to other sort of inefficiencies. For, clearly, the waiting time for swaps between ground level chain and its more tempered counterparts will increase together with the overall number of temperatures. This is because if at each step we try to swap only two chains then the probability of proposing to swap the chain with original target with others decreases. It implies that the best candidate for maximal temperature is to take the smallest temperature at which the distribution becomes unimodal. Notice, that if more swaps occur between the tempered levels, then we do not really solve the problem of poor mixing at the ground level.

There are two possible solutions to the above-mentioned problem contingent upon our knowledge on π. If the number of modes is known in advance, one could simply observe the created trajectories and consider additional temperature levels whenever the algorithm did not arrive at finding some modes. Triming the temperature schedule could also be arranged if the distribution was known to be unimodal at more than one level. In general, the number of modes is unknown and such procedures would not be feasible, and the only safe solution remaining is to overestimate the number of levels and adaptively decrease it. That is precisely the approach we follow. In practice this leaves still one parameter to be provided by the user — the initial number of temperatures.

It is easy to construct toy problems with the temperatures extending over too limited a range. In such

cases the *Parallel Tempering* algorithm would not find other modes except the one where it was initialized. In such cases nearly all criterions to test the target multimodality would fail. To decide that on some level target distribution is unimodal or not we will use comparison of covariance matrices of proposal of random walk and of target.

The well known results on scaling of random walks Metropolis (Roberts et al, 1997) for i.i.d targets and further extensions for more general form of target distributions) choice of the optimal covariance of proposal distribution. This choice is to set proposal covariance equal $\frac{2.38^2}{d}\Sigma$, where Σ is the covariance of the stationary distribution, which corresponds to stationary acceptance probability equal 0.234. There are two popular ways of adaptive schemes which tries to approximate this optimal choice. The first is to estimate on-line the stationary covariance by $\Sigma^{[n]}$ and in next step use proposal with a covariance $\frac{2.38^2}{d}\Sigma^{[n]}$ (Haario et al, 2001). A standard way to define $\Sigma^{[n]}$ follows a two step procedure[4]:

– Get the estimator of the target's mean by

$$\mu^{[n]} = (1-\gamma^{[n]})\mu^{[n-1]} + \gamma^{[n]} X^{[n-1]}. \quad (8)$$

– Get the estimator of the target's covariance matrix

$$\Sigma^{[n]} = (1-\gamma^{[n]})\Sigma^{[n-1]} + \\ + \gamma^{[n]}(X^{[n-1]} - \mu^{[n]})(X^{[n-1]} - \mu^{[n]})^{\mathrm{t}}. \quad (9)$$

In (8) and (9)

$\gamma^{[n]}$ denotes a deterministic step-size[5].

The second one is to use proposal with covariance of the form $\exp(2\theta^{[n]})\Sigma^{[n]}$, where $\theta^{[n]}$ is adjusted in order to get 0.234 acceptance rate (Andrieu and Thoms, 2008; Roberts and Rosenthal, 2009) by the following a Robbins-Monro type procedure

$$\theta^{[n]} = \theta^{[n-1]} + \gamma^{[n]}(\eta^{[n-1]} - 0.234), \quad (10)$$

where $\eta^{[n-1]}$ denotes the acceptance probability of the random walk phase.

In the case of unimodal distribution both approaches are approximately equivalent. In the multimodal case it is obvious that $\exp(2\theta^{[n]})$ will be significantly lower than $\frac{2.38^2}{d}$. It is motivated by the need to shrink the covariance of proposal so that it mimics the covariance of the target distribution truncated to current mode, i.e. where the chain currently is at.

The number of temperature levels is adjusted according to the following formula

$$L^{[n]} = \min\left\{\ell \in \{1,\ldots,L^{[n-1]}\} : e^{\theta_\ell^{[n]}} \geq \frac{2.38}{\sqrt{d}}\right\}, \quad (11)$$

[4] For simplicity we do omit the chain level index, ℓ.

[5] A usual choice is $\gamma^{[n]} \propto n^{-\alpha}$, where $\alpha \in (0.5, 1)$.

if the condition is satisfied for at least one ℓ and $L^{[n]} = L^{[n-1]}$, otherwise.

It seems plausible to froze adaptation during initial burn-in period until the estimators of covariances $\Sigma_\ell^{[n]}$ and the scaling factors $e^{\theta_\ell^{[n]}}$ stabilise.

Remark 3 A similar criterion to that presented in (11) uses a different threshold, namely

$$\min\left\{\frac{2.38}{\sqrt{d}}, \max_{\ell=1,\ldots,L^{[n-1]}}\left\{(1-\epsilon)e^{\theta_\ell^{[n]}}\right\}\right\}.$$

The above condition is more lenient than the one presented in (11), still being optimal when applied to the Gaussian target. In general $\frac{2.38^2}{d}\Sigma$ might be more than achieved using the optimal scaling, this being equal to $\frac{2.38^2}{I\,d}$ in case of a target density in product form, where I denotes the Fisher information of the shift parameter (Roberts et al, 1997). This stems from the Rao-Crammer lower bound, $I^{-1} \leq \Sigma$.

5 Algorithm

We shall now pass to detailed description of our algorithm.

In the beginning we choose the initial positionings of all the chains arbitrary. We then apply iteratively the basic step of the *Adaptive Parallel Tempering* as described in Algorithm 1. Each step depends on the outputs of the previous step through the positionings of chains in the state space, current estimator of the target covariance matrix and target expected value, the covariance scaling factor, and the details of the temperature scheme such as precise values of temperatures and their overall number. Also, a sequence of descending numbers $\gamma^{[n]}$ is being used.

The basic step consists of five phases.

In the first phase a simple random walk is carried out independently on every chain. The proposal for next step is drawn from the multivariate normal distribution with the input covariance matrix scaled by the $\exp(2\theta^{[n]})$. The proposals are being accepted or rejected following the standard procedure assuring the reversibility of the corresponding kernel.

In the second phase we perform random walk adaptation scheme. It consists in updating estimates of both the expected value and the covariance matrix of the target distribution. Also the proposal covariance scaling factor is being updated so that to make it dependent of the distance between the proposal acceptance calculated in the random walk phase and the theoretical optimal probability of acceptance, mentioned in Section 4. Every update is done using a sequence of weights $\gamma^{[n+1]}$

constructed so as to give less and less attention to the sampled values while the algorithm proceeds. It is also worth mentioning that no explicit calculations are done using full covariance matrices. Instead, we operate on their Cholesky decompositions and update them using the so called *rank-one update*, whose cost is quadratic with matrix dimension. Also, it is easier to draw the proposal in the *Random Walk Phase* having the covariance matrix already decomposed.

In the third phase random swaps between different chains are being performed according the rules described in Section 3.

In the fourth phase the temperature scheme gets adapted. It's being done so as to have the ratio of differences between adjacent temperatures reflect the difference between the last steps probability of accepting the swap between the two levels and the theoretical value of 0.234. The above-mentioned probability is calculated as if the swaps were drawn uniformly from the set of only neighbouring temperature levels alone. The rationale behind using the artificial acceptance rates stems from our requirement to make that decision independent of the number of temperature levels. In particular, the EE strategy is dependent upon the number of temperature levels (see Section 6).

In the fifth phase we adapt the number of temperatures by cutting all the temperature levels with their corresponding square-root of the variance scaling factor being larger than the value of the optimal covariance of the proposal kernel.

We now pass to questions regarding the theoretical underpinnings of the algorithm described above. The stability issues regarding parameters $\Sigma^{[n]}$, $\beta^{[n]}$, and $\mu^{[n]}$ are hard to analyse in general; thus, we proceed with examination thereof while restricting ourselves to the case where projections of parameters can be found altogether on some compact space. The precise description of the above-mentioned projections can be found in (Miasojedow et al, 2013a).

To obtain the Law of Large Numbers we need to replace (Miasojedow et al, 2013b, Lemma 20) by an analogous one which covers other swap strategies.

Lemma 1 *If swap strategy satisfies $p_{ij}(x) = p_{ij}(T_{ij}(x))$ for all x and all $i < j$, and p_{ij} is Lipschitz continuous as a function of inverse temperatures β. Then kernel S is Lipschitz continuous in the V-norm with respect to the inverse temperature β.*

Proof To explicitly denotes dependence upon β we will use notation S_β, $\alpha_{ij}(x,\beta)$, and $p_{ij}(x,\beta)$ for S, α_{ij} and p_{ij} respectively. By definition of kernel S_β for any mea-

Algorithm 1 One step of *Adaptive Parallel Tempering*

Input $X^{[n]}, \Sigma^{[n]}, \theta^{[n]}, \mu^{[n]}, L^{[n]}, \beta^{[n]}$

Random Walk Phase

for $\ell := 1$ to $L^{[n]}$ **do**
  $G \sim \mathcal{N}(0, \Sigma_l^{[n]})$
  $Y_\ell := X_\ell^{[n]} + \exp(\theta_\ell^{[n]}) \times G$
  $\log \eta_\ell^{[n]} := 0 \wedge \beta_l^{[n]} \Big(\log\big(\pi(Y_\ell)\big) - \log\big(\pi(X_l^{[n]})\big) \Big)$
  if $\mathcal{U}(0,1) < \eta_\ell^{[n]}$ **then**
    $\tilde{X}_\ell^{[n]} := Y_\ell$
  else
    $\tilde{X}_\ell^{[n]} := X_\ell^{[n]}$
  end if
end for

Random Walk Adaptation

for $\ell := 1$ to $L^{[n]}$ **do**

$$\mu_\ell^{[n+1]} := (1-\gamma^{[n+1]})\mu_\ell^{[n]} + \gamma^{[n+1]}\tilde{X}_\ell^{[n]}$$
$$\Sigma_\ell^{[n+1]} := (1-\gamma^{[n+1]})\Sigma_\ell^{[n]} + \gamma^{[n+1]}(\tilde{X}_\ell^{[n]} - \mu_\ell^{[n]})(\tilde{X}_\ell^{[n]} - \mu_\ell^{[n]})^{\mathrm{t}}$$
$$\theta_\ell^{[n+1]} := \theta_\ell^{[n]} + \gamma^{[n+1]}(\eta_\ell^{[n]} - 0.234)$$

end for

Random Swap Phase

Compute probabilities $\mathbf{p} := \left\{ p_{ij}(\tilde{X}_\ell^{[n]}) \right\}_{1 \le i < j \le L^{[n]}}$
Sample $(i,j) \sim \mathbf{p}$
$\alpha := \frac{p_{ij}\left(T_{ij}(\tilde{X}^{[n]})\right)}{p_{ij}(\tilde{X}^{[n]})} \left(\frac{\pi(\tilde{X}_i^{[n]})}{\pi(\tilde{X}_j^{[n]})} \right)^{\beta_j^{[n]} - \beta_i^{[n]}} \wedge 1$
if $\mathcal{U}(0,1) < \alpha_{ij}$ **then**
  $X^{[n+1]} := T_{ij}(\tilde{X}^{[n]})$
else
  $X^{[n+1]} := \tilde{X}^{[n]}$
end if

Temperature Scheme Adaptation

$T_1^{[n+1]} = 1$
for $\ell := 1$ to $L^{[n]} - 1$ **do**

$$\xi_\ell := \left(\frac{\pi(X_\ell^{[n+1]})}{\pi(X_{\ell+1}^{[n+1]})} \right)^{\beta_{\ell+1}^{[n]} - \beta_\ell^{[n]}} \wedge 1$$
$$\log(T_{\ell+1}^{[n+1]} - T_\ell^{[n+1]}) := \log(T_{\ell+1}^{[n]} - T_\ell^{[n]}) + \gamma^{[n+1]}(\xi_\ell - 0.234)$$

end for
$\beta^{[n+1]} := \frac{1}{T^{[n+1]}}$

Number of Temperatures Adaptation

if $n > N_0$ **then**
  $L^{[n+1]} := \min\left\{ \ell \in \{1, \dots, L^{[n]}\} : e^{\theta_\ell^{[n+1]}} \ge \frac{2.38}{\sqrt{d}} \right\}$
else
  $L^{[n+1]} := L^{[n]}$
end if

surable function g and any β, β' we have

$$
\begin{aligned}
&|S_\beta g(x) - S_{\beta'} g(x)| \\
&= \Big| \sum_{i<j} \left(p_{ij}(x,\beta)\alpha_{ij}(x,\beta) - p_{ij}(x,\beta')\alpha_{ij}(x,\beta')\right) g(T_{ij}(x)) \\
&\quad + \sum_{i<j} \left(p_{ij}(x,\beta')\alpha_{ij}(x,\beta') - p_{ij}(x,\beta)\alpha_{ij}(x,\beta)\right) g(x) \Big| \\
&\leq \sum_{i<j} \Big| \left(p_{ij}(x,\beta)\alpha_{ij}(x,\beta) - p_{ij}(x,\beta')\alpha_{ij}(x,\beta')\right) \Big| \bar{g}_{ij}(x),
\end{aligned}
$$

where $\bar{g}_{ij}(x) = \max\{|g(x)|, |g(T_{ij})(x)|\}$. Since $p_{ij}(x,\beta) \leq 1$ and $\alpha_{ij}(x,\beta) \leq 1$, by triangle inequality we get

$$
\begin{aligned}
&\Big| \left(p_{ij}(x,\beta)\alpha_{ij}(x,\beta) - p_{ij}(x,\beta')\alpha_{ij}(x,\beta')\right) \Big| \\
&\qquad \leq |p_{ij}(x,\beta) - p_{ij}(x,\beta')| + |\alpha_{ij}(x,\beta) - \alpha_{ij}(x,\beta')|
\end{aligned}
$$

Terms whit differences of p_{ij} are Lipshitz by assumptions the Lipshitz conditions for terms with α_{ij} follows by the same arguments as in proof of Lemma 20 from (Miasojedow et al, 2013b)

Theorem 1 *Under assumptions of Lemma 1 and regularity conditions as in Theorem 5 of (Miasojedow et al, 2013a) for any function f such that $\sup_x |f(x)\pi(x)^\tau| < \infty$, where $\tau < \frac{\beta_{min}}{2}$ with β_{min} being the lower bound of the compact set of allowed temperatures, it is true that*

$$
\frac{1}{n}\sum_{k=1}^{n} f(X_1^{[k]}) \to \pi f \quad a.s.\ ,
$$

where $X_1^{[n]}$ is a first component of APT algorithm with fixed number of the temperatures levels.

Proof The proof can be carried out using the same arguments as the one to Theorem 5 of (Miasojedow et al, 2013a), with only one minor modification consisting in explicit use of the Lipschitz continuity of swap kernel given by Lemma 1.

Remark 4 It is true the number of temperatures is a non-increasing function of the iteration and hence almost every trajectory has asymptotically a constant number of levels. This fact suggests that the LLN might be satisfied with adaptation of the temperatures number.

6 Numerical Experiments

We have carried out computer simulations to test the functioning of our algorithms. In this section we propose two case studies of application of our algorithms together with a deepened analysis or the results obtained with them. Both of the provided examples are inherently multimodal.

Our first object of interest was testing the efficiency of our algorithm on an example with a controlled number of modes. An adequate toy-example used in literature as benchmark for testing MCMC algorithms comes from the article of Liang and Wong (2001). There, a mixture of 20 normal peaks with equal weights is being considered. The variances of these modes are small enough to keep the the spikes well separated, as might be inspected in Figure 1.

Fig. 1: Twenty Gaussian Peaks Toy Example

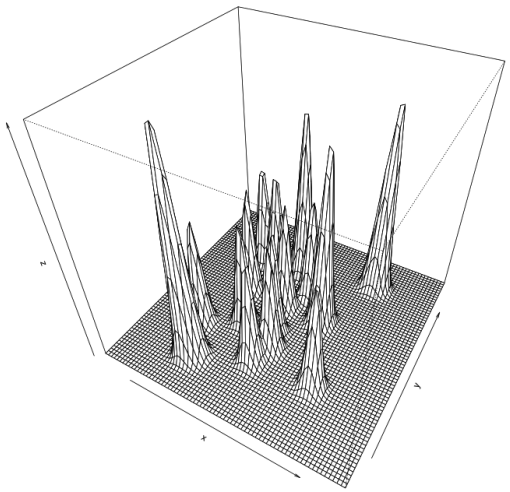

Observe that some peaks form more dense probability clusters that are still separated by large bands of almost zero probability regions.

Table 1: *Root Mean Square Error* of estimators of the two first moments.

RMSE	$\mathbb{E}X_1$	$\mathbb{E}X_2$	$\mathbb{E}X_1^2$	$\mathbb{E}X_2^2$
L=2	0.87	1.16	13.71	11.25
L=3	0.36	0.50	9.24	4.96
L=4	0.33	0.41	9.25	4.32
L=5	0.33	0.41	9.05	4.16

Calculated for 2D Gaussian mixtures toy target using Equi-Energy type swap strategy for a different number of temperature levels.

We have performed simulations (500 runs) to check the quality of the algorithms. Since the actual first and

Table 2: The efficiency of different algorithms in discovering modes of a 2D target distribution.

		$L=3$			$L=4$			$L=5$		
Simulations' Feature	$L=2$	EE	AL	RA	EE	AL	RA	EE	AL	RA
No missing modes (%)	62.8	99.8	99.6	99.8	100	99.2	99.4	100	99.4	98.8
b Mean Absolute Error	0.59	0.32	0.34	0.35	0.3	0.36	0.37	0.29	0.38	0.4

Summary of the results of the *Parallel Tempering* algorithm applied to the 2D Gaussian mixtures toy target with number of temperature levels equal to $L \in \{2,3,4,5\}$. Three different swap strategies were tested: Equi Energy type moves (EE), swapping only two adjacent levels (AL), and swapping between all levels at random (RA). The *Mean Absolute Error* of time spent in each mode is approximated by calculating $\frac{1}{20}\sum_{i=1}^{20}\frac{|t_i-0.05|}{0.05}$, t_i being the proportion of time spent in i^{th} mode.

second moments of this distribution are readily obtainable by direct calculation, we could evaluate the *Root Mean Square Error* of the moments estimators, as gathered in Table 1, and compare ratios thereof while using different swap strategies, see Table 3. Every simulation consisted of 2,500 iterations of burn-in followed by 5,000 steps of the main algorithm.

Table 3: Ratio of RMSE of estimators of the two first moments

Temperatures N°	Ratio	$\mathbb{E}X_1$	$\mathbb{E}X_2$	$\mathbb{E}X_1^2$	$\mathbb{E}X_2^2$
$L=3$	AL/EE	1.04	0.96	1.01	0.97
	RA/EE	1.04	0.97	1.05	1.00
$L=4$	AL/EE	1.13	1.27	1.00	1.21
	RA/EE	1.23	1.20	1.05	1.18
$L=5$	AL/EE	1.30	1.31	1.06	1.27
	RA/EE	1.24	1.30	1.06	1.30

Calculated for 2D Gaussian mixtures target for different number of temperatures and different choice of swap strategy. Confront Table 2 for swap strategy encodings.

We have also modified the upper mentioned example by considering the product of that distribution with six independent uniform distributions. The final distribution is composed therefore of 20 peaks that are being stretched to additional dimensions. The purpose of that experiment was to test how our algorithms cope with the difficulty of exploring a highly multidimensional space. Results of numerical experiments (100 runs) are contained in Table 4. Each run was preliminated by 5,000 steps of burn-in, followed by 10,000 steps of the proper algorithm.

To see how an exemplary trajectory manages its search for probability modes confront Figures 2 and 3.

Fig. 2: Twenty Gaussian Peaks 2D Simulation

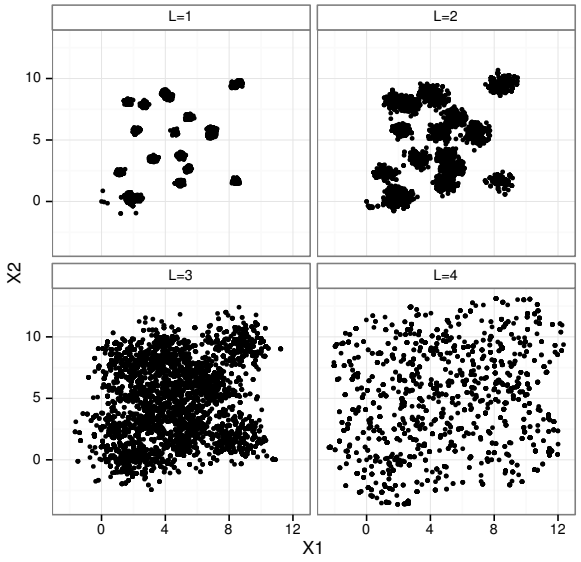

Here we plot one individual trajectory. Each subplot represents the natural projection of that trajectory on the state space containing the original multimodal distributions and their transforms. The simulation lasted 7500 iterations, preceded by 2500 iterations of burn-in. The criterion for temperature scheme reduction was not set on until the burn-in period has finished. When turned on, it reduced the problem to three levels out of original four.

It can be clearly seen that not at all temperature levels can the algorithm easily travel between different modes.

It was of interest to us how well do the algorithms perform when dealing with the inherent multimodiality of the upper-mentioned problems. In case of two first problems mentioned above we do know the number of modes a priori and checking whether an algorithm managed to explore a particular mode could be dealt with use of clusterings of the state space into regions representing particular modes. Tables 2 and 4 do present

Table 4: The efficiency of different algorithms in discovering modes of a 8D target distribution.

	$L=3$			$L=5$			$L=7$			$L=9$		
Simulations' Feature	EE	AL	RA	EE	AL	RA	EE	AL	RA	EE	AL	RA
No missing modes (%)	0	0	0	17.3	16.3	5.1	17.3	11.2	3.06	19.4	6.12	0
Average number of missing modes	16	16	16	1.6	1.6	2.7	1.8	2.2	3.4	1.6	2.4	5.2
Mean Absolute Error	1.7	1.7	1.7	0.78	0.79	0.89	0.79	0.86	0.93	0.8	0.85	1.1

Summary of the results of the *Parallel Tempering* algorithm applied to the 8D Gaussian mixtures toy target with number of temperature levels equal to $L \in \{3, 5, 7, 9\}$. Confront Table 2 for swap strategies' encodings and definition of the *Mean Absolute Error*.

Fig. 3: Twenty Gaussian Cylinders 8D Simulation

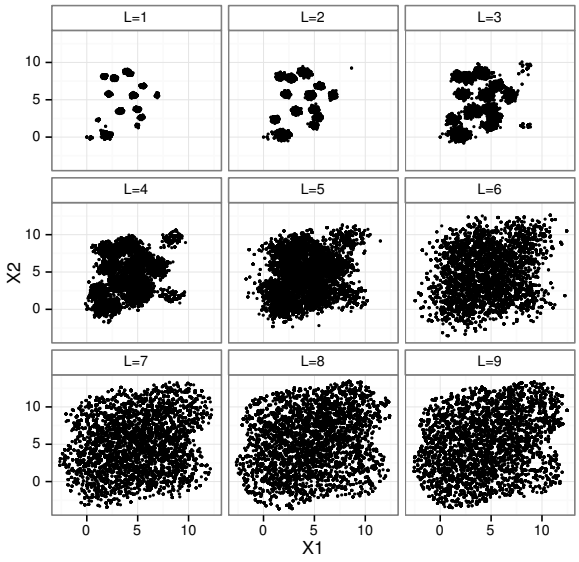

Here we plot one individual trajectory. Each subplot represents the natural projection of that trajectory on the state space containing the original multimodial distributions and their transforms.
The simulation lasted 15,000 iterations, preceded by 5,000 iterations of burn-in. The criterion for temperature scheme reduction was not set on until the burn-in period has finished. When turned on, it reduced the temperature scheme usually to four or five levels. When the criterion has been turned on after 7,500 iterations, it indicated always that one should consider five levels.

the efficiency of estimates as measured by the expected number missed modes during one run of the algorithm and also the frequency of case when the algorithms does not miss any mode.

We have run the algorithms in two modes. In first mode we have neglected temperature scheme reduction (but maintained the temperature adaptation as such). These simulations were carried out to provide insight into how the initial number of temperatures affects different strategies results. The results for the 2D toy example are gathered in Tables 2-3. It can be generally observed that the *Equi Energy* (EE) strategy slightly outperforms the state-independent swap strategies, the difference getting bigger with more temperature levels being taken into account. Clearly this can be attributed to the fact that a state-dependent swap is more immune to the quadratic growth in the number of possible swaps: choosing a swap uniformly from all potential swaps renders exchanges between the base level and other levels less likely in overall. On the other hand, choosing only adjacent levels (AL) seems to be only slightly worse.

We have therefore tumbled upon an important rationale for temperature scheme adaptation as such: if the probability of proposed swaps is not concentrated on a relatively small subset, then it is less likely for the algorithm to convey information about a new mode discovery to the base chain. Observe however, that one cannot *a priori* underestimate the proper number of temperatures; for in such case the tempered chains could be still exploring poorly separated probability modes of the transformated target distribution. For instance, in Figure 4 one can explicitly see that there is a huge difference between the errors in first two moments estimates for the 8D toy example when one uses 3 temperature levels instead of 4 or more. Similarly, in Table 2 one notices that AL and RA strategies do perform worse with more temperature levels considered, for on average the number of modes that they miss slightly grows. It is completely different in case of the state-dependent EE strategy that works actually better with the growing number of temperatures. The supremacy of this state-dependent strategy can be clearly seen in Table 3 that presents direct comparison of the *Root Mean Square Errors* of other strategies when compared to that of the EE strategy. Tables 4-5 show similar results for the 8D toy example.

Table 5: Ratio of the *Root Mean Square Error* of estimators of the two first moments.

	L=3		L=4		L=5		L=6		L=7		L=8		L=9	
Moment	AL/EE	RA/EE	AL/EE	RA/EE	AL/EE	RA/EE	AL/EE	RA/EE	AL/EE	RA/EE	AL/EE	RA/EE	AL/EE	RA/EE
$\mathbb{E}X_1$	1.00	1.00	1.08	1.02	0.97	1.03	1.03	0.94	1.33	1.38	0.95	1.14	1.03	1.16
$\mathbb{E}X_2$	1.01	1.02	1.07	1.00	1.15	0.98	0.97	1.06	1.12	1.11	0.95	1.35	1.13	1.51
$\mathbb{E}X_1^2$	1.00	1.00	1.02	1.01	0.93	1.05	0.96	0.91	1.04	1.02	0.96	1.08	1.01	1.08
$\mathbb{E}X_2^2$	1.01	1.01	0.95	1.00	1.09	0.97	1.04	1.11	1.15	1.24	1.01	1.43	1.05	1.55

Calculated for 8D Gaussian mixtures target for different number of temperatures with different swap strategies, at different temperature levels. Confront Table 2 for swap strategy encodings.

Table 6 shows also another interesting phenomenon: the overall percentage of accepted swaps grows with the number of temperature levels when applying the EE strategy. It suggests that the algorithm performs well in finding points in the state-space with similar probabilities. It also suggests that that truly the temperature adaptation procedure should not be using the original probabilities of acceptances but some other quantities independent of the number of initial temperatures. This provides a rationale for our choice of probabilities that would result from application of any swap proposal kernel that is symmetric (see the description of the fourth phase of the algorithm in Section 5).

Table 6: Average acceptance rate of swap steps.

Strategy	L=3	L=4	L=5	L=6	L=7	L=8	L=9
EE	0.38	0.42	0.45	0.45	0.46	0.46	0.46
RA	0.16	0.12	0.10	0.08	0.07	0.06	0.05

Calculated for 8D Gaussian mixtures target for Equi-Energy type swap strategy and random swap strategy at different temperature levels Confront Table 2 for swap strategy encodings.

Fig. 4: Root mean square error of estimators of first and second moments

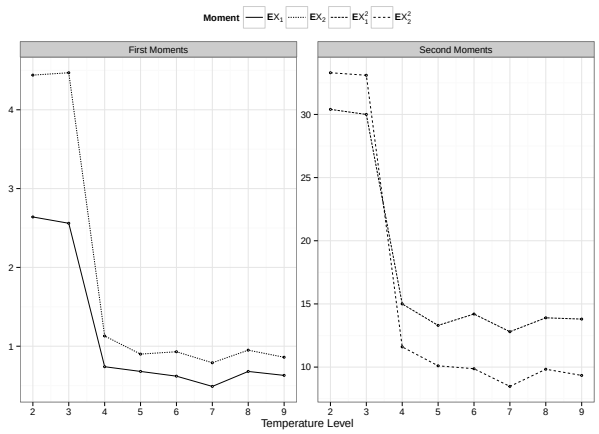

Calculated for 8D Gaussian mixtures target for different number of temperatures with Equi-Energy type swap strategy at different temperature levels. Apparently there exists a threshold value of the initial number of temperatures that renders the estimates much more accurate. Not knowing it a priori, we are bound to overshot the true value to obtain high quality estimates. This is also a rationale behind any automated temperature reduction algorithm.

The general message from the above analysis is clearly that there exists a threshold number of temperature levels that significantly ameliorates the *Parallel Tempering* algorithm's performance and that the state-dependent strategy might partially solve problems resulting in unlikely swaps with the base level chain. We shall now describe the second phase of our numerical experiments with the toy examples, namely the verification of how well does our temperature scheme reduction criterion works in practice.

To this end we have carried out yet another series of experiments (100 runs). In the 2D toy example (7,500 steps, out of which 2,500 burn-in) the ultimate number of temperature levels in every simulation reached three. In the 8D example the result depended on the burn-in period: if out of 15,000 steps the burn-in amounted to 7,500 then the algorithm always reached five levels of temperatures. When given a shorter burn-in period in 85% of cases it reached five levels, otherwise descending even to four. It seems therefore that the algorithm approaches levels that could have been intuitively chosen and, what is important, in a more conservative way: it simply at worse slightly overestimates the actual level that might have been chosen by visual inspection of the "scree plots" such as the ones depicted in Figure 4.

To show that the algorithms developed in this article can be of use in applications, we have tested them on the problem of the estimation of linear model coefficients. The formulation of the problem comes from Park and Casella (2008) and is known in the literature under the name of *Bayesian Bridge* regression. It consists in performing Bayesian inference on the parameters of a standard linear model

$$\boldsymbol{y} = \mu \mathbf{1}_n + \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where one postulates that $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}_n, \Sigma)$, resulting in

$$\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{X}\boldsymbol{\beta}, \sigma \mathbb{I}_n),$$

and additionally one poses a shrinkage prior on the $\boldsymbol{\beta}$ coefficients

$$\pi(\boldsymbol{\beta}) = \prod_{j=1}^{p} \frac{\lambda}{2} e^{-\lambda |\beta_j|^q}.$$

Here, p is the number of different features gathered in the data set X, and n is the overall number of sample points. Finally, one poses a scale invariant prior on the model variance $\sigma \sim \pi(\sigma) = \frac{1}{\sigma^2}$. There is an explicit relation between the *maximum a posteriori* estimate of parameters in the above mentioned model and the solution to the problem of coefficients estimates provided by the Lasso algorithm (Tibshirani, 1996). The *a priori* independence of the $\boldsymbol{\beta}$ from σ^2 might render the $\boldsymbol{\beta}$ *a posteriori* distribution multimodial, as in the case studied herein - following Park and Casella (2008) we have run the algorithms on the DIABETES dataset taken from Efron et al (2004).

In this example the starting point for the algorithms (namely - a pair $(\boldsymbol{\beta}^{[0]}, \sigma^{[0]})$) was set to be a simple OLS estimate for the coefficients $\hat{\boldsymbol{\beta}}$ together with the resulting mean residual sum of squares being a proxy for real variance, $\hat{\sigma}$. We have carried out one hundred simulations each consisting of 20,000 iterations of burn-in and 100,000 iterations of the actual algorithm, for different initial number of temperatures. The results of our simulations are represented in Table 7. While performing the simulations it was observed that a simple random walk could not leave the neighbourhood of $\mathbf{0}_n$ even though most of the probability mass lays elsewhere, as can be investigated from the plot for the first temperature level ($L = 1$) in Figure 5.

Results on the final number of reached temperatures are gathered in Figure 6 and demonstrate that its inclusion in the algorithm gives rather conservative results in terms of temperature levels reduction, which might be considered safe an option.

Table 7: Posterior Means of selected Parameters of the *Bayesian Bridge* regression

	β_3	β_9	$\log(\sigma)$
RWM	0.228 (0.64)	0.165 (0.390)	8.69 (0.003)
EE	723 (2.60)	0.240 (0.150)	8.3 (0.0031)
AL	724 (2.60)	0.199 (0.050)	8.3 (0.0027)
RA	723 (4.40)	0.141 (0.067)	8.3 (0.0028)

Summary of the results of the *Bayesian Bridge* regression. It can be seen that the simple random walk Metropolis (RWM) strategy does not succeed in exploring all the modes. When it comes to other strategies, no significant differences can be spotted.

Fig. 5: Posterior Distribtuion of Coefficients in Bridge Regression

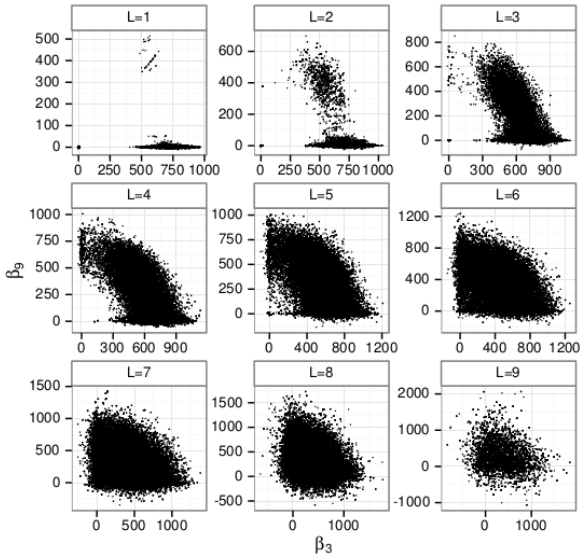

Above, projections of one simulated chain on the subspace spanned by two parameters in the Bayesian Lasso model. Visual inspection of base temperature ($L = 1$) reveals the inherent multimodiality of the *a posteriori* distribution. The simulations were performed on nine different temperature levels. One notices visually that at temperature ranging from the 6^{th} and 7^{th} level the modes are not easily recognisable, which implies that the corresponding chains no longer experience local behaviour. Also our criterion (cerified every 1000 steps of the algorithm) did cut off levels higher than 7.

7 Conclusions

In this article it was demonstrated that the *Parallel Tempering* algorithm's efficiency is highly dependent upon the number of initial temperature levels; for underestimating that number leads to poor quality esti-

Fig. 6: Distribution of Ultimate Number of Temperature Levels in the *Bayesian Bridge* Regression

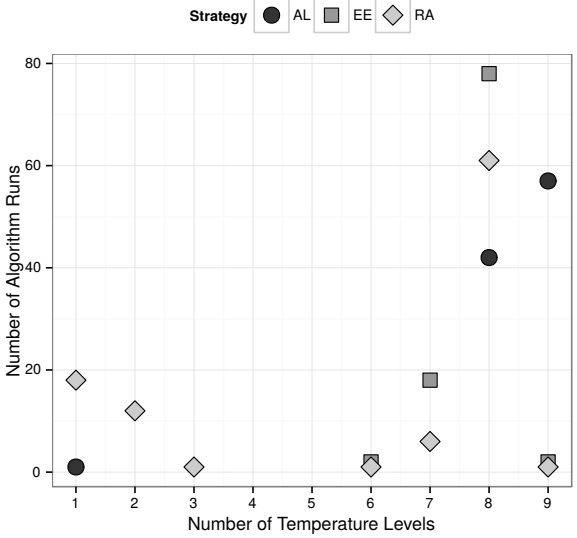

The above plot counts how many times different numbers of temperature were reached when using the temperature number adaptation and different swap strategies. The simulations lasted 100K iterations, preceded with 20K steps of burn-in. It is interesting to observe that the EE strategy never reached a number smaller than six and usually promoted a conservative choice of eight levels of temperature. Confront that with Figure 3.

mates resulting from neglecting some of the modes of the distribution of interest. On the other hand, overestimation of the number of temperatures does not significantly improve these estimates but highly increases the computational cost. Finally, if in that case one uses a naive state-independent strategy, then the overall efficiency of the algorithm drops due to longer waiting times for significant swaps, i.e. ones involving the base level temperature.

We have developed an algorithm with adaptable temperature scheme. Using it, one can overshot the number of temperature levels needed for performing good quality simulation and the procedure will automatically reduce some of the redundant levels. The results of our simulation support our claim of the method correctness.

We have also proposed a novel for state dependent strategies that allow promotion of swaps based on various properties of the current state, e.g. equi-energy type moves, big-jumps promotion, and similar. We have shown that this framework is susceptible to analytical analysis based on already known results. Specifically, we have concentrated here on evaluating the equi-energy type strategy and thoroughly tested it. The results show that it is more efficient than the standard state-independent swap strategy, since the overall acceptance of a global moves is higher, the difference getting larger with the initial number of temperatures. Similar conclusions can be drawn on the behaviour of the *Root Mean Square Error*.

When it comes to future research, we plan to prepare an **R**-package containing all the above methods.

All the scripts used in simulations are readily obtainable on request.

Acknowledgements We are grateful to Matti Vihola for useful discussions.

References

Andrieu C, Thoms J (2008) A tutorial on adaptive MCMC. Statistical Computing 18(4):343–373

Atchadé YF, Rosenthal JS (2005) On adaptive Markov chain Monte Carlo algorithms. Bernoulli 11(5):815–828

Atchadé YF, Roberts GO, Rosenthal JS (2011) Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. Statistics and Computing 21(4):555–568

Baragatti M, Grimaud A, Pommeret D (2013) Parallel tempering with equi-energy moves. Statistics and Computing 23(3):323–339

Efron B, Hastie T, Johnstone I, Tibshirani R (2004) Least angle regression. The Annals of Statistics (32):407–499

Gilks WR, Richardson S, Spiegelhalter DJ (1998) Markov Chain Monte Carlo in Practice. Chapman & Hall/CRC, Boca Raton, Florida

Haario H, Saksman E, Tamminen J (2001) An adaptive Metropolis algorithm. Bernoulli 7(2):223–242

Hansmann UHE (1997) Parallel tempering algorithm for conformational studies of biological molecules. Chemical Physics Letters 281(1-3):140–150

Kofke D (2002) On the acceptance probability of replica-exchange Monte Carlo trials. Journal of Chemical Physics 117(15):6911–1914

Kone A, Kofke D (2005) Selection of temperature intervals for parallel-tempering simulations. Journal of Chemical Physics 122

Kou SC, Zhou Q, Wong WH (2006) Equi-energy sampler with applications in statistical inference and statistical mechanics. The Annals of Statistics 34(4):1581–1619

Liang F, Wong WH (2001) Evolutionary Monte Carlo for protein folding simulations. Journal of Chemical Physics 115(7)

Marinari E, Parisi G (1992) Simulated tempering: a new Monte Carlo scheme. Europhys Lett 19(6):451–458

Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E (1953) Equations of state calculations by fast computing machines. Journal of Chemical Physics 21(6):1087–1092

Miasojedow B, Moulines E, Vihola M (2013a) An adaptive parallel tempering algorithm. Journal of Computational and Graphical Statistics 22(3):649–664

Miasojedow B, Moulines E, Vihola M (2013b) Appendix to "An adaptive parallel tempering algorithm". Journal of Computational and Graphical Statistics

Park T, Casella G (2008) Bayesian lasso. Journal of the American Statistical Association 103(482):681–686

Robert CP, Casella G (1999) Monte Carlo Statistical Methods. Springer-Verlag, New York

Roberts GO, Rosenthal JS (2009) Examples of adaptive MCMC. Journal of Computational and Graphical Statistics 18(2):349–367

Roberts GO, Rosenthal JS (2012) Minimising MCMC variance via diffusion limits, with an application to simulated tempering. Tech. rep., URL `http://probability.ca/jeff/research.html`

Roberts GO, Gelman A, Gilks WR (1997) Weak convergence and optimal scaling of random walk Metropolis algorithms. Annals of Applied Probability 7(1):110–120

Schreck A, Fort G, Moulines E (2013) Adaptive equi-energy sampler: Convergence and illustration. ACM Trans Model Comput Simul 23(1):5:1–5:27

Swendsen RH, Wang JS (1986) Replica Monte Carlo simulation of spin-glasses. Physical Review Letters 57(21):2607–2609

Tibshirani R (1996) Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B (58):267–288

Woodard DB, Schmidler SC, Huber M (2009) Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. Annals of Applied Probability 19(2):617–640